Tracking the onset of productive determiner+noun combinations in English-learners

Ruthe Foushee[a]
Dan Byrne[a]
Raquel G. Alhama[c]
Allyson Ettinger[b]
Afra Alishahi[c]
Susan Goldin-Meadow[a]

[a]Department of Psychology
University of Chicago
5848 S University Ave
Chicago, IL 60637, USA

[b]Department of Linguistics
University of Chicago
115 E 58th St
Chicago, IL 60637, USA

[c]Department of Cognitive Science and Artificial Intelligence
Tilburg University
5000 LE Tilburg
The Netherlands

Author Note

Ruthe Foushee, https://orcid.org/0000-0002-8361-1725, @ruthefoushee
Dan Byrne, @UChicagoLDP
Raquel G. Alhama, https://orcid.org/0000-0002-8573-6716, rgalhama
Allyson Ettinger, https://orcid.org/0000-0001-7701-7888, @AllysonEttinger
Afra Alishahi, @afraalishahi
Susan Goldin-Meadow, https://orcid.org/0000-0001-8536-2673

Correspondence should be addressed to:

Dr. Ruthe Foushee (foushee@uchicago.edu)
Department of Psychology
5848 S University Ave
Chicago, IL 60637, USA

Acknowledgments

Research reported in this publication was supported by the Eunice Kennedy Shriver National Institute of Child Health & Human Development of the National Institutes of Health under Award Number P01HD040605; by a grant from the Successful Pathways from School to Work initiative of the University of Chicago, funded by the Hymen Milgrom Supporting Organization; by a grant from the Institute of Education Sciences, U.S. Department of Education, through grant R305A190467 to the University of Chicago; and by a grant from the Spencer Foundation. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health, the Institute of Education Sciences, or the U.S. Department of Education.

Abstract

Previous work shows that *determiner+noun* productivity is not present at the earliest moments of language acquisition (Meylan et al., 2017). Here we develop a new onset measure of productivity that can be used with less data, and we apply it to longitudinal observations of 64 English-learners interacting with their parents at home. We confirm the original finding and show, for the first time, that when productivity onsets, it applies not just to *a/the* + noun combinations, but to the entire *class* of determiners. We also find, again for the first time, that the onset of productivity is *not* predicted by parental linguistic input. Our findings do not constitute evidence for or against preformed abstract linguistic categories. However, our new onset measure, which allows us to assess productivity at the earliest stages of language-learning, sets the stage for addressing this question in future computational work.

Keywords: language acquisition, linguistic productivity, grammatical category learning, syntactic development, corpus methods

Introduction

Having heard "a pimwit," English-speakers know immediately that "the pimwit" is possible, even if they haven't heard the phrase before. At what age do children do the same? Researchers from diverse theoretical perspectives agree that this type of productivity provides evidence of abstract syntactic categories (NOUN, DETERMINER). The disagreement lies in *when* young language-learners display this productivity, which bears on (but does not determine) whether children come to language-learning with abstract categories. For researchers arguing for early-available abstract syntactic categories, the phonetic forms *a* and *the* in English are learned as members of the pre-formed DETERMINER category. For researchers arguing for gradually-constructed abstract syntactic categories, the DETERMINER category itself must be generalized from repeated experiences with *a* and *the*. Interestingly, researchers from opposing theoretical perspectives have largely used the same data to support their arguments.

Early contributions to this debate analyzed syntactic errors in children's productions to determine how often children produce *a* or *the* in the wrong position within a noun phrase (e.g., *big the dog*), or stacked in sequence (e.g., *the a dog*), or without a noun at all. Analyzing utterances from six 2-year-old children, Valian (1986) found only one determiner error and argued that, to be this accurate in the earliest moments of language-learning, children must have abstract syntactic categories guiding their acquisition. Subsequent research focused on an extension of this claim. If children have abstract noun and determiner categories early in language acquisition, then the same noun should be used with different determiners (i.e., a child should say *a dog* and *the dog)*. If syntactic categories are gradually constructed, a child who hears *a dog* may not immediately understand that *the dog* is also possible. On this reasoning,

when a child begins to use the same noun with both *a* and *the,* the child may be providing evidence that they possess abstract determiner and noun categories.

In seeking a quantitative measure of productivity that could be applied to spontaneous child productions, Pine and Lieven (1997) developed an overlap score. The overlap in question is between the set of nouns used with *a* and the set of nouns used with *the*. Specifically, a child who produced *a dog, a sock, the dog,* and *the plant* would receive an overlap score of 1/3, or 33%, having produced one noun type out of three (*dog, sock, plant*) with both *a* and *the* (*a dog*, *the dog*). Pine and Lieven applied this metric to the first 400 multiword utterances recorded in a study of 12 children from their first to their third birthdays. In their initial analysis, only one of the 11 children in their sample had an overlap score significantly different from zero, providing no evidence that children's determiner category was abstract at this age.

Citing issues with assumptions underlying Pine and Lieven's (1977) overlap score, Yang (2013) proposed an alternative. Certain nouns in English are more likely to appear with one of the two determiners (e.g., "*a* sky" is less attested than "*the* sky;" "*the* wish" is less attested than "*a* wish"). As a result, some nouns are not likely to occur with both determiners in spontaneous discourse. Yang (2013) computed an *expected* overlap score (the overlap score that children would receive if they used language like adults, that is, a score that takes into account how likely a noun in adult speech is to occur with both determiners). When applied to a sufficiently large corpus, children's *expected* and *observed* overlap scores ought not differ *if* children have an abstract determiner category. In addition, children's observed scores ought not differ from caregivers' observed scores. Yang (2013) used a dense corpus of transcribed naturalistic interactions between caregivers and children to confirm both hypotheses.

The overlap measure has not been used to assess overlap scores on the same child at different moments in development and thus has not given us a picture of individual children's developmental trajectories of *determiner+noun* productions. In other words, the measure has not been used to address the *when* question with which we began. Meylan and colleagues (2017) used a Bayesian model to address this question. Under their model, a child's determiner productions for each of their nouns are guided by two information sources—(a) direct experience and (b) productive knowledge. The strength of each source's contribution to the child's productions is determined by individual weighting parameters. Meylan et al. found low levels of productivity initially and higher levels later in development in 26 children. They argued that the findings were consistent with the view that children lack abstract grammatical categories at the outset of language learning, but rapidly begin to develop the categories on the basis of their linguistic input. Note, however, that just because children do not *display* productivity when they begin producing determiners does not mean that they come to language-learning without an abstract determiner category—the children may possess an abstract determiner category but need time to discover what that category looks like in the language they are learning. Nevertheless, pinpointing *when* a child becomes productive with respect to determiners and nouns is important in its own right, and also puts empirical boundaries on our theories of abstract category development.

Our goal here is to introduce a new measure of the onset of children's productive use of *determiner+noun* combinations, one that has several advantages over previously used measures. Yang's (2013) overlap measure has not been used to track individual trajectories. Meylan et al.'s (2017) model has been used to track individual trajectories, but the researchers found that the model works best when applied to a great deal of data from each child. Here we use a type of

overlap score and credit a child with productive use of *determiner+noun* combinations when the child uses both *a* and *the* with at least two different nouns (i.e., *a baby, the baby, a birdie, the birdie*). Our onset criterion does not require that a child talk a great deal (a child could theoretically demonstrate productivity with just 4 noun phrases). This feature allows us to measure productivity in individual children using data that are much less dense than Meylan et al.'s analysis requires. We apply our measure to a naturalistic corpus of child language that begins at 14 months—10 months before Meylan et al. (2017) saw the rapid increase in their measure. Our sample also includes many more children, and extends over a longer period of time, than all previous samples, allowing us to characterize the developmental trajectory of *determiner+noun* combinations in many children.

Our first goal enables us to address two additional goals. First, we examine onset productivity of other determiners (possessives and demonstratives) to determine whether they onset at the same time as *a/the+noun* combinations. If so, we will have identified the onset of a *determiner class*. Second, we explore whether the talk each child hears from their parents predicts the time in development when the child begins to productively combine determiners with nouns, a question that has not been previously explored.

Method

Corpus

The data for this study come from the Language Development Project corpus, which follows 64 typically developing, monolingual, English-speaking children from the Greater Chicagoland Area (Goldin-Meadow et al., 2014); 64 children and their primary caregivers were video-recorded engaging in spontaneous interactions in their homes for twelve 90-minute visits

(M=11.3, SD = 1.8, sessions, range 4–12 sessions), beginning from when the children were 14 months to 58 months. The resulting corpus of caregiver-child interactions contains over 1 million transcribed utterances (n = 646,685 for primary caregivers and n = 368,884 for children), and approximately 1,000 hours of videos. For a complete description of the corpus, see Goldin-Meadow et al. (2014).

Data Selection, Preprocessing, and Criterion for Onset Productivity

Both the primary caregiver's and child's utterances were lemmatized, stripped of extraneous punctuation, and all instances of capitalization were removed. All utterances tagged as reading were excluded. Following this preprocessing, *determiner+noun* pairs were identified by leveraging the open-source natural language processing library SpaCy (Honnibal & Montani, 2017) to assign parts-of-speech tags to every word within each utterance. Using these parts-of-speech tags, we extracted determiner-(adjective)-noun sequences containing determiners from three categories: indefinite/definite (*a/an, the*), possessive (*my, your, his, her(s), our(s) their(s)*), and demonstrative (*this, that, these, those*). Following previous literature (Meylan et al., 2017; Pine & Lieven, 1997; Pine & Martindale, 1996; Pine et al., 2013; Yang, 2013; Yang & Valian, 2020), we focus first on combinations of indefinite/definite determiners with nouns. Our preprocessing pipeline and analysis code is archived at https://osf.io/s2jnm/?view_only=ca2d57aee759426ba1c531a64bc982f0.

We assume that a child has productive use of *a/the+noun* combinations when the child uses both *a* and *the* with the same noun and (to insure that the combination is not a fluke) does so with at least two different nouns (i.e., *a car, the car, a bottle, the bottle*). In order to meet this criterion, the child has to produce at least two instances of two different nouns within a session.

Results

Using *A* and *The* Productively with Nouns

We begin by examining when children first produced *a* and *the* in our corpus. The children first produced *a* or *the* between 14 months (their first visit; $n = 6$) and 38 months ($n = 1$); mean age = 21.80 (SD = 4.83) months. Many children produced *a* before *the* ($n = 33$), although almost as many used both *a* and *the* on their first session ($n = 27$). Only 4 children produced *the* before *a.*

Figure 1a shows the number of children who first met our criterion for onset productivity at each age. Of our sample of 64 children, 63 met the criterion for productivity for combining both *a* and *the* with the same noun, and doing this twice, within our 12 observation sessions. The number of sessions between a child's first determiner and meeting our definition of onset productivity varied from 0 sessions ($n = 4$) to 7 sessions ($n= 1$); mean number of sessions between first appearance and productivity = 2.29, SD = 1.36.

Although requiring substantially less data for detecting productivity than previous methods, our measure nevertheless requires that a child produce at least two different nouns twice. Not surprisingly, the median number of words children in our sample produced across sessions was, in fact, correlated with their onset of productivity (Kendall's rank correlation $\tau = -0.25$, $p=.007$)—children who talk more have more opportunities to meet our criterion than children who talk less. However, those opportunities did not determine when most of the children first produced *a* and *the* with the same noun and did this with at least two different nouns. We examined the observation sessions *prior* to the session at which each child met our criterion, and asked whether the child had produced at least two different nouns twice during this

period. We found that 53 children (84%) had produced two nouns two times on at least one session preceding that child's onset of productivity. In other words, the children had produced enough noun phrases that they *could* have met our onset productivity criterion—but they didn't, suggesting that we had captured the time period when children were first productive. For the remaining 10 children (16%), their first productive session was the first time that the child met the enabling conditions for productivity.

Figure 1

(a) Number of Children Classified According to the Age at Which They Met the Criterion for A/The+Noun *Onset Productivity (b) Median Number of Noun Types Children Combined with* A *and* The *at Each Observation Session, Classified According to the Age at Which the Child First Met the Criterion for Onset Productivity*

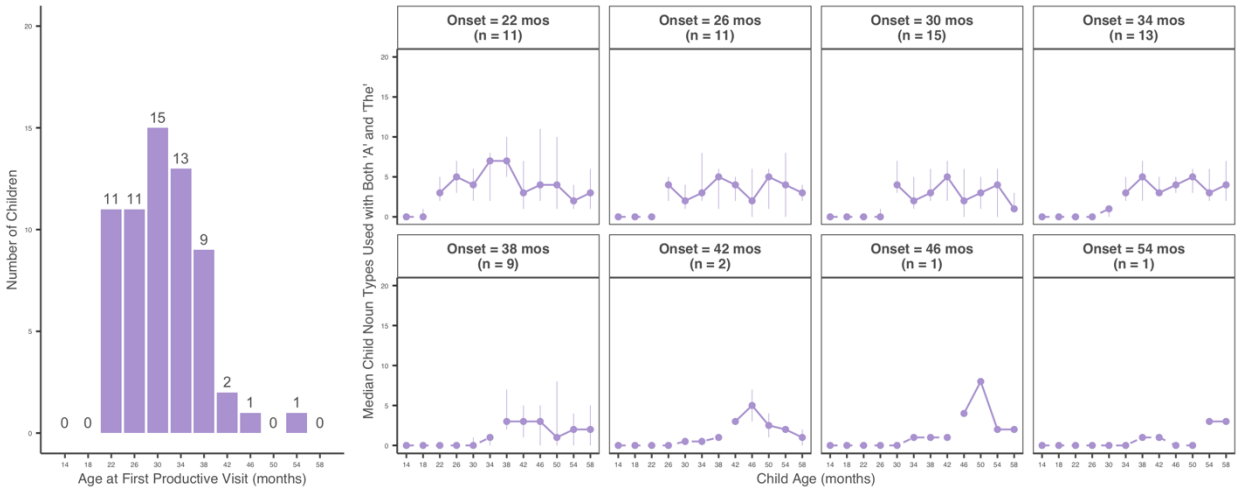

Note. One child did not meet the criterion during our observations and is excluded from both plots.

Figure 1b presents the median number of different nouns that children produced with both *a* and *the* at each age; the children are categorized according to the age at which they first produced two different nouns with both *a* and *the.* Note that, after having met the criterion, children in each of the groups produced roughly the same number of different nouns with both *a* and *the,* no matter when they first achieved productivity. The exceptional groups (the groups who first displayed productivity at 42, 46 or 54 months) contained very few children (2, 2, and 1, respectively). The majority of children (52 of 63, 83%) met our productivity criterion on at least half of the sessions following their onset; 22 children reached criterion on every subsequent session.

Applying the Overlap Measure to our Sample

To assess how our productivity measure compares to previously used measures, we used Yang's (2013) procedure to determine the overlap between the set of nouns used with *a,* and the set of nouns used with *the*, for each child in our sample. Because the 12 children in the Manchester corpus (Theakston et al., 2001) produced an average of 52 unique nouns (SD=7.3) in their earliest stages, Yang (personal communication) estimates that a child would need to produce at least 50 unique nouns in order for the overlap measure to be valid. We therefore looked at the first session when each of the children in our sample produced 50 unique nouns (mean age=35.6, SD=8.1), and calculated an overlap score for that session. We found that, in all cases, there were no significant differences between how likely the child was to display productivity (the expected score) and how often the child *did* display productivity (the observed score); see *Supplementary Materials*. In other words, our children behaved like English-speakers with respect to *a/the+noun* combinations as soon as they produced 50 unique nouns.

However, by the time they produced 50 unique nouns, all but one of the children in our sample were productive according to our measure. The advantage of our measure is that it can be used with children who produce few noun types. In fact, the mean number of noun types children produced during their first productive session (according to our measure) was 33.3 (SD=18.6, range from 8 to 84), considerably fewer than the number needed for the overlap score. Our measure has thus allowed us to explore the very earliest stages of productivity in individual children. Note, however, that the lack of *determiner+noun* productivity in the earliest sessions found in our measure (and in Meylan et al.'s, 2017) cannot be taken as evidence for the absence of preformed categories, a point to which we return in the Discussion.

Evidence for an Abstract Determiner Category

Previous studies have examined *a* and *the* as an instance of the determiner category. But *possessives* and *demonstratives* are also part of the determiner category in English and are routinely positioned before nouns. We therefore examined children's use of possessives [*my, mine, your(s), our(s), his, her(s), their(s)*] and demonstratives [*this*, *that*, *these*, those] during this time period.

Children produced their first possessives between 14 months (n = 1) and 34 months (n = 1); mean age = 24.30 (SD=4.2) months, and their first demonstratives between 14 months (n = 1) and 38 months (n = 1); mean age = 26.00 (SD=4.73) months. Using the same criterion for onset productivity that we used for *a* and *the*, we found that children first combined two possessives with the same noun, and also with a second noun, between 30 months (n = 1) and 58 months (n = 2); mean age = 45.40 (SD=8.13) months. The children first combined two demonstratives with the same noun, and also with a second noun, between 30 months (n = 1) and 58 months (n = 4)*;*

mean age = 46.30 (SD=7.79) months. Thirty-three children used possessives productively and 43 used demonstratives productively during our observation sessions. In terms of continued productivity, 4 met our productivity criterion on at least half of the sessions after onset of productivity for possessives, and 5 did so for demonstratives.

The productivity onsets for the three types of determiners are roughly comparable, suggesting that the children's early productions may have been guided by a single, abstract determiner category. To further explore this possibility, we went back to the data and treated all three types of determiners (definites/indefinites, possessives, demonstratives) as members of a single category, and examined the onset of productivity and continuity for this inclusive determiner category. We asked when children first began producing two determiners (*a, the*, *my, mine, your(s), our(s), his, her(s), their(s), this*, *that*, *these*, *those*) with the same noun and did this twice.

Figure 2a presents the number of children who first met our criterion for onset productivity of *determiner+noun* combinations at each age. All 64 children in our sample met the criterion for productivity in *determiner+noun* combinations within our observation sessions. The distribution of onset ages in Figures 1a and 2a are comparable, a bell-shaped curve around approximately the same mean age (M=30.9, SD=6.6 for *a/the+noun* combinations, Fig.1a; M=28.2, SD=5.3 for *all determiner+noun* combinations, Fig.2a). Onset of productivity did not change for 35 of the 64 children under this new analysis. For those whose onset did change, 20 began productivity one session earlier (at M=31.2, SD = 6.24), 5 began two sessions earlier (at M=25.2, SD=3.35), and 4 began three sessions earlier (at M=28.0, SD=5.2). Importantly, all but four children (60 out of 64, 94%) continuously met the criterion for *determiner+noun* productivity after having produced their first *determiner+noun* combination. Figure 2b presents

the median number of distinct nouns that children produced with two different determiners at each session, categorized according to the first session at which the child met the criterion for productivity.

Figure 2

(a) Number of Children Classified According to the Age at Which They First Met the Criterion for Determiner+Noun *Onset Productivity (b) Median Number of Noun Types Children Combined with two Different Determiners at Each Observation Session, Classified According to the Age at Which the Child Met the Criterion for Onset Productivity.*

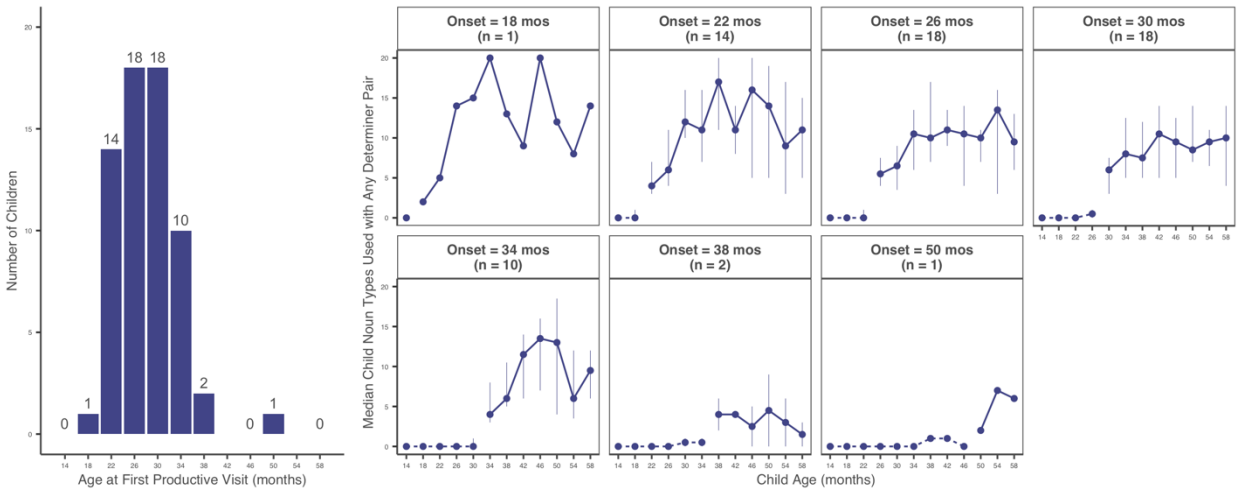

Productivity in the Speech Children Receive from their Caregivers

We applied the same criterion for productivity to the speech of the children's primary caregivers, first for combining nouns with *a* and *the* (Figure 3) and then for combining nouns with any of the three types of determiners (Figure 4). Not surprisingly, 52 out of 64 caregivers produced at least two different nouns with both *a* and *the* at the first observation session (see

Figure 3a), and 62 out of 64 caregivers produced at least two different nouns with two exemplars of the *determiner category* at the first observation session (see Figure 4a). In terms of continuity, 62 of the caregivers met our criterion for combining *a/the* with different nouns on all sessions following their onset; all 64 of the caregivers met our criterion for combining two determiners with different nouns at all sessions after their onset.

Figure 3

(a) Number of Caregivers Classified According to When They Met the Criterion for A/The+Noun *Productivity (b) Median Number of Noun Types Caregivers Combined with A and The at Each Observation Session, Classified According to the Age at Which the Caregiver's Child Met the Criterion for Productivity*

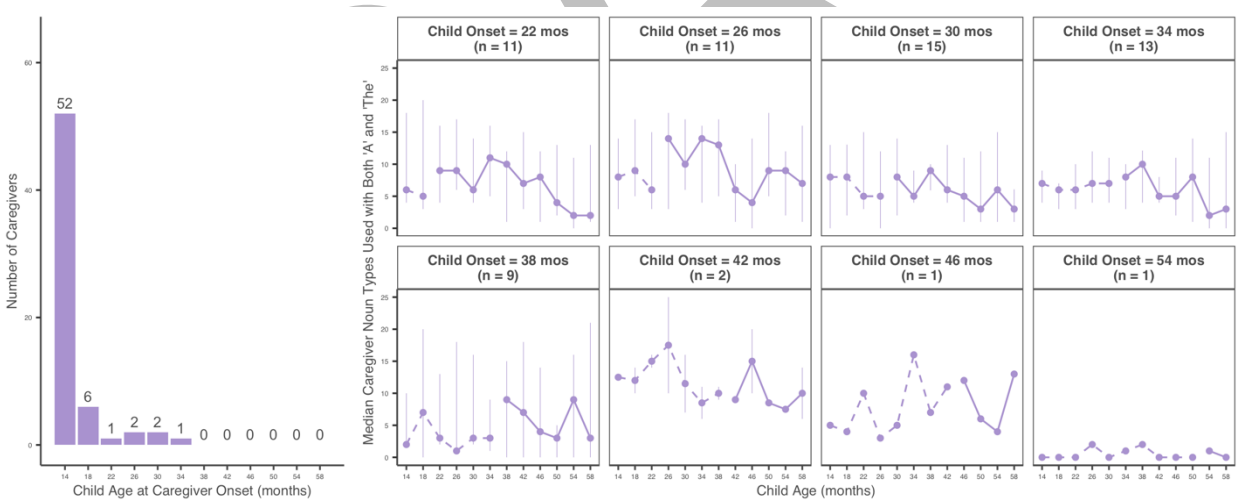

We also calculated the median number of noun types caregivers combined with *a/the* (Figure 3b) or with any of the three types of determiners (Figure 4b) at each observation session; caregivers were classified according to the age at which the caregiver's child met the respective

criterion for productivity. Note that the number of different nouns that the caregiver combined with *a/the,* or with any of the three determiners, did *not* change before or after the children in any of the groups achieved onset productivity. This point suggests that the *timing* of caregiver input did not determine when individual children achieved productivity. Nevertheless, the *amount* and *type* of caregiver input may have had an impact on when children first displayed productivity in combining nouns with *a/the* or with any of the determiners. We turn to this question in the next section.

Figure 4

(a) Number of Caregivers Classified According to When They Met the Criterion for Determiner+Noun *Productivity (b) Median Number of Noun Types Caregivers Combined with two Different Determiners at Each Observation Session, Classified According to the Age at Which the Caregiver's Child Met the Criterion for Productivity*

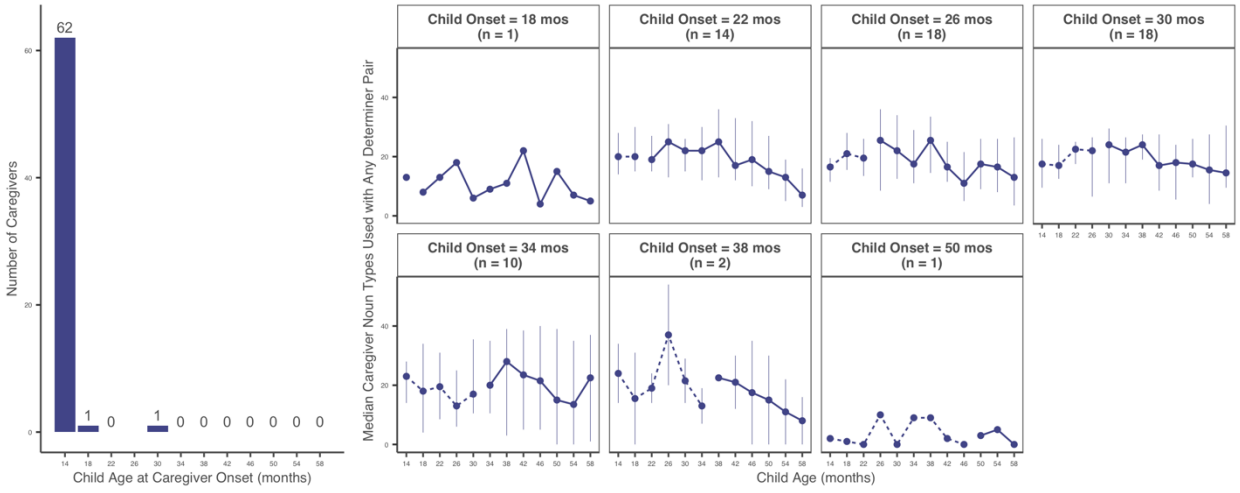

Does Type of Caregiver Input Predict Onset of Child Productivity?

We assessed the total number of words (tokens), the total number of unique nouns (types), and the total number of *determiner+noun* combinations that each caregiver produced at sessions one and two because none of the children had met our criterion for onset productivity during these sessions. Caregivers produced a mean number of total words (3723, SD=1457, range=811—8814), of total unique nouns (72, SD=29, range=11—147) and of *determiner+noun* combinations (154, SD=79, range=23—446) during the first two sessions.[1] Note that the variability across caregivers was quite large for each measure, allowing us to ask whether these measures of caregiver input could account for the onset of productive *a/the+noun* and *determiner+noun* combinations.

We fit linear models to children's onset of productive *a/the+noun* combinations (Table 1) and to the onset of productive *determiner+noun* combinations (Table 2), using each of these three measures of caregiver input as the sole predictor of onset. None of the measures of caregiver input was related to the timing of children's onset productivity.

Table 1

Linear Models Predicting Child Onset of Productive A/The+Noun *Combinations from Caregiver Input*

	Dependent variable:			
	Onset of determiner+noun productivity			
Caregiver Measures	(1)	(2)	(3)	(4)
Intercept	31.95***	33.51***	33.16***	32.71***

[1] Mean caregiver-produced noun types before session 2 predict mean child-produced noun types after session 2 (b=0.14 [0.04, 0.25], F(1, 56)=7.07, p=.01), and mean caregiver-produced determiner-noun combinations before session 2 predict mean child-produced determiner-noun combinations after session 2 (b=0.40 [0.31, 0.50], F(1, 56)=72.39, p<.001).

	(27.95, 35.94)	(29.08, 37.95)	(29.52, 36.80)	(29.52, 36.80)
Total number of words (tokens)	-0.0003 (-0.001, 0.001)			
Total unique nouns (types)		-0.03 (-0.08, 0.02)		
Total number of *determiner+noun* phrases			-0.01 (-0.02, 0.004)	
Total number of different nouns combined with both *a* and *the*				-0.23 (-0.48, 0.03)
Observations	63	63	63	63
R2	0.01	0.03	0.03	0.05
Adj. R2	-0.01	0.01	0.02	0.03
Residual Std. Error (*df* = 61)	6.68	6.61	6.60	6.53
F Statistic (*df* = 1; 61)	0.34	1.63	1.96	3.06

**p<.05, **p<.01, ***p<.001*

Table 2

Linear Models Predicting Child Onset of Productive Determiner+Noun *Combinations from Caregiver Input*

	Dependent variable:			
Caregiver Measures	*Onset of determiner+noun productivity*			
	(1)	(2)	(3)	(4)
Intercept	29.54*** (26.36, 32.73)	29.91*** (26.34, 33.47)	29.76*** (26.83, 32.69)	29.49*** (26.74, 32.23)
Total number of words (tokens)	-0.0004 (-0.001, 0.001)			
Total unique nouns (types)		-0.019 (-0.06, 0.02)		

Total number of *determiner+ noun* phrases			-0.007 (-0.02 0.005)	
Total number of different nouns combined with distinct determiners				-0.065 (-0.19 0.06)
Observations	63	63	63	63
R2	0.01	0.02	0.02	0.02
Adj. R2	-0.003	0.001	0.006	0.002
Residual Std. Error (df=61)	5.32	5.32	5.30	5.31
F Statistic (df = 1; 61)	0.84	1.04	1.40	1.12

**p<.05, **p<.01, ***p<.001*

Discussion

We have explored the onset of *determiner+noun* productivity in the early period of language-learning using a measure that has a number of advantages over previously used measures. The measure is intuitive, easily applied to individual children, and needs less dense data than other measures. Accordingly, we were able to observe many more children over a longer period of time than in previous studies. Our data confirm, in 64 children, that *determiner+noun* productivity is not present at the outset, but begins relatively early in development (Meylan et al., 2017)—the modal age of productivity onset was 30 months in our data, and ranged from 22 to 54 months. It's possible that our measure underestimates productivity (i.e., children might be productive from the beginning, but our measure may be too insensitive to detect it). However, most children met the enabling conditions for our measure (at least two instances of a noun, for two nouns) *prior* to meeting our criterion for productivity. In other words, the enabling conditions of our measure were not the limiting factor on detecting

productivity. Moreover, once children displayed productive use of *a/the* + *noun* combinations, they continued to be productive, suggesting that they had indeed made a linguistic leap (see Figure 1b).

Our second finding is that, when children first begin to productively combine nouns with *a* and *the,* they also combine nouns with different types of determiners (indefinite/definite, possessive, demonstrative). Although previous work has described the determiner category as containing different types of specifiers (e.g., Zamparelli, 2000), the argument about when children first acquire a determiner category has centered around their production of *a/the* + *noun combinations* (Meylan et al., 2017; Pine et al., 2013; Pine & Lieven, 1997; Pine & Martindale, 1996; Yang, 2013; Yang & Valian, 2020). Our data are the first to suggest that a broad abstract determiner category is present when children first display productivity in their *a/the* + noun combinations.

Finally, ours is the first study to examine the impact that caregiver input has on the *onset* of child *determiner+noun* productivity. We found that the number of productive instances of these combinations that caregivers produced did not vary before and after the onset of productivity in the child. Moreover, we were unable to predict the onset of child productivity from any of our measures of caregiver input despite considerable variability across caregivers in these measures. This finding highlights the importance of child-centered individual differences in the onset of productivity.

Although we have pinpointed the onset of *determiner+noun* productivity in a large sample of children, our findings (like Meylan et al.'s, 2017) do not tell us whether children come to language-learning equipped with an abstract determiner category. Our findings do, however, set the stage for exploring this question using computational modeling, which we are currently

pursuing using our data as a starting point (Alhama et al., *in prep*). Having characterized at a fine-grained level the timing associated with development of *determiner+noun* productivity in different individual children, and the properties of caregiver input for each child, we have generated a rich empirical target for computational simulation. We therefore have the opportunity not only to design computational models that embody different hypotheses about the language-learner and the learning mechanisms, but also to explore the extent to which these hypotheses interact with observed caregiver input to produce the individual developmental trajectories described here.

There is reason to believe that our computational modeling will find that children need little caregiver input to produce *determiner+noun* combinations. The data come from homesigners. Homesigners are profoundly deaf children whose hearing losses prevent them from acquiring a spoken language, and whose hearing parents have not exposed them to a signed language. Despite their lack of linguistic input, these children communicate and use gestures, called *homesigns,* to do so. Their homesigns display many, although not all, of the properties of natural language (Goldin-Meadow, 2003, 2020). One of the properties that *is* found in homesign is *determiner+noun* combinations (Hunsicker & Goldin-Meadow, 2012). Importantly, the homesigners' hearing parents do *not* display *determiner+noun* combinations in the spontaneous gestures that they produce when they talk to their deaf children (Flaherty et al., 2021), indicating that the children do not have a model for this category either from a conventional language or from spontaneous gesture. The homesign observations suggest that children come to language-learning prepared to create a determiner category if they are not receiving linguistic input. Our findings here suggest that, if children do receive linguistic input, they will search for a

determiner category in that input and achieve productivity in *determiner+noun* combinations around 30 months.

In sum, we confirmed previous findings that children do not display *determiner+noun* productivity at the earliest moments of language acquisition. We also found that, when it begins, productivity applies not just to *a/the* + noun combinations, but to the entire *class* of determiners. Finally, we found that the onset of productivity is not easily predicted by caregiver linguistic input. Our findings do not constitute evidence for or against preformed abstract linguistic categories, but our paradigm sets the stage for addressing this question in future computational work.

References

Flaherty, M., Hunsicker, D., & Goldin-Meadow, S. (2021). Structural biases that children bring to language learning: A cross-cultural look at gestural input to homesign. *Cognition*, *211*, 104608. https://doi.org/10.1016/j.cognition.2021.104608

Goldin-Meadow, S. (2003). *The Resilience of Language*. Psychology Press. https://doi.org/10.4324/9780203943267

Goldin-Meadow, S. (2020). Discovering the Biases Children Bring to Language Learning. *Child Development Perspectives*, *14*(4), 195–201. https://doi.org/10.1111/cdep.12379

Goldin-Meadow, S., Levine, S. C., Hedges, L. V, Huttenlocher, J., Raudenbush, S. W., & Small, S. L. (2014). New evidence about language and cognitive development based on a longitudinal study: Hypotheses for intervention. In *American Psychologist* (Vol. 69, Issue 6, pp. 588–599). American Psychological Association. https://doi.org/10.1037/a0036886

Honnibal, M., & Montani, I. (2017). *spaCy 2: Natural language understanding with Bloom*

embeddings; convolutional neural networks and incremental parsing.

Hunsicker, D., & Goldin-Meadow, S. (2012). Hierarchical structure in a self-created communication system: Building nominal constituents in homesign. *Language*, *88*(4), 732–763. https://doi.org/10.1353/lan.2012.0092

Meylan, S. C., Frank, M. C., Roy, B. C., & Levy, R. (2017). The Emergence of an Abstract Grammatical Category in Children's Early Speech. *Psychological Science*, *28*(2), 181–192. https://doi.org/10.1177/0956797616677753

Pine, J. M., Freudenthal, D., Krajewski, G., & Gobet, F. (2013). Do young children have adult-like syntactic categories? Zipf's law and the case of the determiner. *Cognition*, *127*(3), 345–360. https://doi.org/10.1016/j.cognition.2013.02.006

Pine, J. M., & Lieven, E. V. M. (1997). Slot and frame patterns and the development of the determiner category. *Applied Psycholinguistics*, *18*(2), 123–138. https://doi.org/10.1017/s0142716400009930

Pine, J. M., & Martindale, H. (1996). Syntactic categories in the speech of young children: the case of the determiner. *Journal of Child Language*, *23*(2), 369–395. https://doi.org/10.1017/S0305000900008849

Theakston, A. L., Lieven, E. V. M., Pine, J. M., & Rowland, C. F. (2001). The role of performance limitations in the acquisition of verb-argument structure: an alternative account. *Journal of Child Language*, *28*(1), 127–152. https://doi.org/DOI: 10.1017/S0305000900004608

Yang, C. (2013). Ontogeny and phylogeny of language. *Proceedings of the National Academy of Sciences*, *110*(16), 6324 LP – 6327. https://doi.org/10.1073/pnas.1216803110

Yang, C., & Valian, V. (2020). *Determining the Abstractness of Determiners*.

Zamparelli, R. (2000). *Layers in the Determiner Phrase*. New York: Garland Publishing. https://doi.org/9780815337690